arXiv:1410.7383v6 [stat.ML] 18 Sep 2015

A Ternary Non-Commutative Latent Factor Model for Scalable Three-Way Real Tensor Completion

G. Baruch, Yahoo Labs, guy.baruch@icloud.com

First online version: 26 October 2014

Abstract

Motivated by large-scale Collaborative-Filtering applications, we present a Non-Commuting Latent Factor (NCLF) tensor-completion approach for modeling three-way arrays, which is diagonal like the standard PARAFAC, but wherein different terms distinguish different kinds of three-way relations of co-clusters, as determined by permutations of latent factors.

The first key component of the algebraic representation is the usage of two non-commutative real trilinear operations as the building blocks of the approximation. These operations are the standard three dimensional triple-product and a trilinear product on a two-dimensional real vector space $\mathbb{C}^{\perp} \subset \mathbb{R}^{2\times 2}$, which is a representation of the real Clifford Algebra $Cl(1,1)$ (a certain Majorana spinor). Both operations are purely ternary in that they cannot be decomposed into two group-operations on the relevant spaces. The second key component of the method is combining these operations using permutation-symmetry preserving linear combinations.

We apply the model to the MovieLens and Fannie Mae datasets, and find that it outperforms the PARAFAC model. We propose some future directions, such as unsupervised-learning.

1 Introduction

Tensor completion of three-way arrays[1] had been used to model three-way interactions in many experimental fields, starting in the 1920s with the chemometrics and psychometrics communities. Kolda and Bader provide an extensive review of tensor factorization literature up to 2009 [KB09]. A shorter but more current review is given by Graesdyck et al. in [GKT13].

This work considers three-way interactions in a "Collaborative Filtering" (CF) context. In the classical CF problem, some quantity of interest M (deterministic or stochastic) depends on two variables of large cardinality $M = M(i,j) = M_{ij}$ where $i = 1,\ldots,I$ and $j = 1,\ldots,J$, which is naturally represented as a matrix. The matrix of known values is typically sparse, and the problem is to estimate the missing values, seeking the best approximation in the L^2 (Froebenius) norm. In the three-way case the quantity of interest depends on three variables and is represented as a cuboid tensor $T = T(i,j,k) = T_{ijk}$. See Figure 1 for an illustration and Section 2 for a more concrete example.

[1] Semantics of the term "tensor" differs between research communities, as elucidated in Section 2 of [dSL08]. We will take "tensor" to be equivalent of "n-way array".

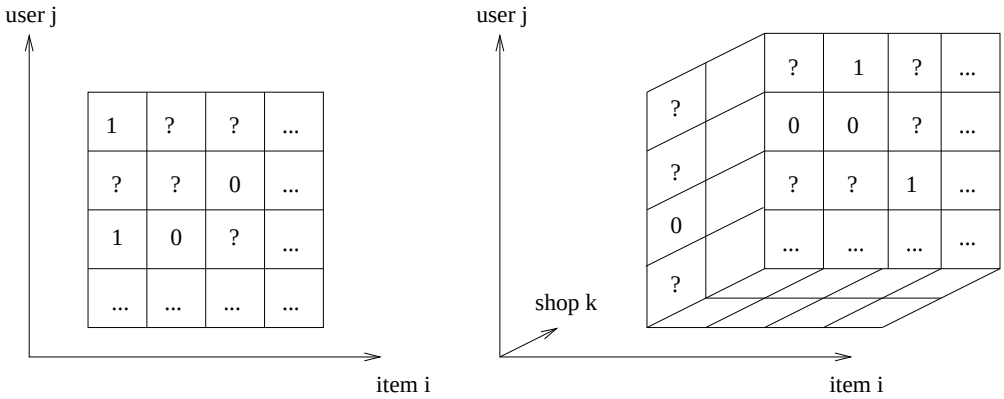

Figure 1: Left: An example of the classical Boolean Collaborative Filtering (CF) problem, wherein a binary response variable (indicating for example a purchase event) is given for each user-item pair, represented as a sparse matrix. Questions marks denote unknown values. The problem is estimating the probability of a purchase events for an unseen pair. Right: The corresponding three-way CF problem we consider, where the response variable depends on a triplet, in this case of user, item and shop, represented as a sparse cuboid.

The two main tensor decompositions used are the CANDECOMP/PARAFAC (CP) model proposed by Hitchcock in 1927 [Hit27b, Hit27a], and the Tucker decomposition proposed by Tucker in 1963 [Tuc66, Tuc63, Tuc64]. In the CP model, a three-way array $T \in \mathbb{R}^{I \times J \times K}$ is approximated by a finite sum of rank-1 tensors

$$T_{ijk}^{\text{CP}} = \sum_{r=1}^{R} U_{ir} V_{jr} W_{kr} + \text{bias terms}, \tag{1}$$

where $U \in \mathbb{R}^{I \times R}$, $V \in \mathbb{R}^{J \times R}$, $W \in \mathbb{R}^{K \times R}$ are called "latent factor matrices", and

$$\text{bias terms} = b_0 + b_{1i} + b_{2j} + b_{3k}.$$

For readers unfamiliar with machine learning terminology, we note that the name "latent factor" stems from an assumption that the data is generated from a fixed distribution governed by variables which are hidden (latent).

In the more general Tucker model the latent factor rows are multiplied by a "core tensor" $C_{R_1 \times R_2 \times R_3}$ of dimensions R_1, R_2, R_3 as

$$T_{ijk}^{\text{TUCKER}} = \sum_{r_1=1}^{R_1} \sum_{r_2=1}^{R_2} \sum_{r_3=1}^{R_3} C_{r_1 r_2 r_3} U_{i r_1} V_{j r_2} W_{k r_3} + \text{bias terms}. \tag{2}$$

The Tucker model is more expressive than the CP model, but its core tensor is typically dense, requiring $\mathcal{O}\left(R^3\right)$ parameters. It is also harder to interpret.

We note that the CP model has the property that its basic building block - the real triple product $u_i v_j w_k$ - does not distinguish between cases wherein the numerical values of the latent factors are permuted, for example between $(u_i, v_j, w_k) = (1, 2, 3)$ and $(u_i, v_j, w_k) = (2, 3, 1)$ (and similarly for other permutations). In other words, for the three-way interactions modeled by CF, a commutative building block is inherently less

expressive than a non-commutative one. Thus, we speculate that three-way relations are better distinguished by a product of non-commuting latent factors than by the (commutative) real multiplication of the CP model. This intuition is expanded in Section 3.

Following this speculation, we propose a hybrid of the CP and the Tucker3 models which is pseudo-diagonal (like the CP), but is built ground-up from trilinear operations of Non-Commuting Latent Factors (NCLF). The general form of the NCLF model is

$$\begin{gathered} T_{ijk}^{\mathrm{NCLF}} = \sum_{\mathrm{sym}=1}^{6} \sum_{r=1}^{R_{\mathrm{sym}}} L_{\mathrm{sym}}\left(U_{ir}^{\mathrm{sym}}, V_{jr}^{\mathrm{sym}}, W_{kr}^{\mathrm{sym}}\right) + \text{bias terms}, \\ U^{\mathrm{sym}} \in \mathbb{V}^{I \times R_{\mathrm{sym}}}, \quad V^{\mathrm{sym}} \in \mathbb{V}^{J \times R_{\mathrm{sym}}}, \quad W^{\mathrm{sym}} \in \mathbb{V}^{K \times R_{\mathrm{sym}}}, \end{gathered} \tag{3}$$

where the subscript "sym" denotes different permutation symmetries of latent factors, $L_{\mathrm{sym}}(\cdot)$ is a real trilinear mapping satisfying this symmetry mode, and $\mathbb{V}$ is a real linear space to be determined.

A well-known problem of unregularized CP models is that approximations of a certain rank may not exist, a situation commonly called "degeneracy", see Section 3.3 of [KB09] and also [CLdA09]. De Silva and Lek Heng Lim show that such degeneracy can be generic, i.e., occurring at a non zero-measure set of inputs [dSL08]. They also prove that degeneracy always co-occurs with the formation of collinear columns of the latent factor matrices, meaning that the set of vectors $U_{:1}, U_{:2}, \ldots, U_{:r}$, where the colon sign denotes a running index, becomes linearly dependent, or almost so. This dependency manifests in very large columns which almost cancel each other. They also note that, while regularization removes non-existence, proximity of the well-posed regularized problem to the ill-posed unregularized problem may still result in catastrophic ill-conditioning.

Much of the effort in lower-dimension tensor factorization have been directed into extending the Singular Value Decomposition (SVD), for example by applying orthogonality constraints on the columns of the latent factor matrices or of the core matrix of the Tucker decomposition - see a review in [Kol01]. Orthogonality of matrix-slices of the Tucker core tensor has been considered by L. de Lathauwer et al., who show that this model retains many properties of the original matrix SVD, therefore naming it the High Order SVD (HOSVD) [dLdMV00]. The core tensor, however, is still dense requiring $\mathcal{O}\left(R^3\right)$ parameters.

When the dimension of the factors is small, orthogonality and collinearity of the latent matrix columns are mutually exclusive, and orthogonality removes degeneracy even for the CP model. For typical "big data" CF problems, however, dimensionality of each factor may be extremely large[2] and so virtually all vector pairs are near-orthogonal. Near-orthogonality is therefore not useful in avoiding collinearity. We note that a standard CP expansion of a finite-rank NCLF model will always have collinear parallel factors. Hence, some degenerate modes may be alleviated by the NCLF model. We leave the question of how much degeneracy is alleviated open[3].

In the completely different setting of particle physics, modeling three-way interactions (in three-quark models) have been shown to be intrinsically related to non-commutativity of the underlying algebras. Kerner proposed using one such algebra in three-color quark models [Ker10], and we shall use such ideas for the

[2] For example, each Yahoo user may receive her own latent row vector, and the number of such users is in the hundreds of millions.

[3] Some examples wherein the CP model becomes degenerate are associated with differential operators, see [dSL08]. The NCLF model directly models CP-degenerate modes associated with first-order finite-difference operators. Therefore, we speculate it removes degeneracy associated with first-order differential operators, but not all the higher-order ones.

algebraic representation used by our model[4].

For the reader familiar with Geometric Algebra we add two notes, which other readers may safely ignore. First, we will use the two dimensional real representation of the Clifford Algebra $Cl(1,1)$, which in Physics is known as one of the flavors of a Majorana spinor. Second, some recent tensor factorization works use Grassman algebras to represent the completely antisymmetric components of the input [KB09, KSV]. In the third order case the standard triple product in $\mathbb{R}^3$, which is the approach we use for this component, is a Grassman Algebra.

The remainder of the paper is as follows. In Section 2 we formulate the specific CF problem we are interested in. In Section 3 we give the motivating intuitions of this work. Specifically, we conjecture that in order to distinguish between three-way relations by a single term, an algebraic representation must be non-commutative. Moreover, it must model, either implicitly or explicitly, different permutation symmetries of the latent factors. Following these intuitions, in Section 4 we construct the NCLF model, which we construct in several steps:

1. In Section 4.1 we recall the decomposition of a generic cubical tensor into its symmetry-preserving components. This decomposition is done via six linear operators.

2. In Section 4.2 we look for and find a non-commutative trilinear mapping μ on a two-dimensional linear subspace $\mathbb{C}^{\perp}$ of $\mathbb{R}^{2\times 2}$, which is the simplest such mapping we could devise. This mapping is the key component of our method, and will be used to construct five of the six symmetry-preserving components of the NCLF model. We denote this space by $\mathbb{C}^{\perp}$ because it is the orthogonal complement of the representation of the Complex field in $\mathbb{R}^{2\times 2}$. The mapping μ is purely ternary, meaning that the space $\mathbb{C}^{\perp}$ is closed under the trilinear operation, but not under the corresponding bilinear one. In other words, $\mathbb{C}^{\perp}$ is a ternary algebra, not a standard (binary) algebra.

3. In Section 4.3 we approximate each of these components by its own trilinear mapping: the completely antisymmetric component is modeled by the standard triple-product in $\mathbb{R}^3$, and approximation of the other components are constructed by applying the symmetrizing operations on the mapping μ. We provide explicit expressions for each of the components.

4. Finally, in Section 4.4 we assemble the full approximation, and apply it to the general cuboid case.

In Section 5, we provide the results of numerical experiments on two publicly available datasets, the MovieLens movie rating dataset and the Fannie Mae Single Family Home Performance dataset. In both cases, the non commutative models outperform the standard CP model. We conclude and discuss future directions in Section 6.

2 A specific Three-way CF problem

The specific problem motivating this paper is that of predicting binary response via three-way CF in supervised learning. In this learning problem, the dependent variable is a Boolean event - like a purchase event,

[4]For the reader unfamiliar with physics we note that the CF problems we consider are entirely different from quantum chromodynamics, so that we can propose much simpler models.

which we denote by $Y \in \{0, 1\}$, and the independent variables belong to three classes of large cardinality, for example users, purchasable items and shopping venues, see Figure 1 on the right.

The learning problem is therefore to estimate the probability of a purchase event $P(Y = 1|i, j, k)$ for an (unseen) triplet $i \in 1, \ldots, I$, $j \in 1, \ldots, J$ and $k \in 1, \ldots, K$. The value of Y_{ijk} for most of the triplets is unknown, making this a tensor completion problem.

We will use a Logistic Regression model, thereby estimating the log-odds of this probability

$$\log\left(\frac{P(Y=1|i,j,k)}{1-P(\cdot)}\right) \approx T_{ijk} \in \mathbb{R}^{I\times J\times K}, \tag{4a}$$

or, equivalently,

$$P(Y=1|i,j,k) \approx \text{Logit}\left(T_{ijk}\right), \qquad \text{Logit}(T_{ijk}) = \frac{1}{1+\exp\left(-T_{ijk}\right)}. \tag{4b}$$

We will be using L^2 (Tikhonov) regularized models and the logistic loss function, so that given a functional form $T(U, V, W, b)$ (like CP, or Tucker3) and data Y_{ijk} (known over a subset of the triplets (i, j, k)), training will consist of the solution of the minimization problem

$$\begin{aligned} U, V, W, b = \arg\min \sum_{\{(i,j,k):Y_{ijk} \text{ known}\}} & -Y_{ijk}\log\left(\text{Logit}(T_{ijk})\right) - (1-Y_{ijk})\log\left(1-\text{Logit}(T_{ijk})\right) \\ & +\lambda\left\|U\right\|^2 + \lambda\left\|V\right\|^2 + \lambda\left\|W\right\|^2, \end{aligned} \tag{5}$$

where the last three terms are the regularization terms, and the parameter λ is the regularization parameter, to be chosen empirically via cross validation.

These four simplifying assumptions - of a supervised learning, binary response problem modeled by logistic regression with L^2 regularization - are applied in order to demonstrate the NCLF model on a concrete problem. Apriori, they only affect the numerical experiments in Section 5. We see no reason why the NCLF model should not apply to other three-way multilinear subspace learning problems.

3 The intuitive motivation

Let us look for the simplest extension to the trilinear CP model, which would still be be diagonal, but would provide a more expressive algebraic representation of a three-way relation between entities, for example between users, purchasable items and venues. Such a representation approximates how a three-way relation affects some measured quantity - for example the odds $P/(1-P)$ of a purchase event - which we take for simplicity to be real. Since we are estimating a real quantity, we consider real trilinear mappings.

Following intuitions from Physics [Ker10], we speculate that non-commutative parallel factors might be more expressive than commutative ones, i.e., that in reality a "green user, blue item red shop" combination is different than a "blue user, green item, red shop" combination, and will lead to a different propensity to purchase. Since the "colors" are arbitrary regions of the latent factor space corresponding to different co-clusters, there is no reason, priory, to assume that a function representing the relation between parameter regions for shops, items and venues be commutative in the latent factors.

Hence, this article raises the following conjecture:

component	$P_{\rm cyc}$	$P_{\rm acyc}$	P_J	P_{12}	P_{23}	P_{31}
S	1	1	NA	1	1	1
A	1	-1	NA	-1	-1	-1
$\mathcal{J}_{31-}$	NA	NA	0	NA	NA	-1
$\mathcal{J}_{31+}$	NA	NA	0	NA	NA	1
$\mathcal{J}_{23-}$	NA	NA	0	NA	-1	NA
$\mathcal{J}_{23+}$	NA	NA	0	NA	1	NA

Table 1: Eigenvalues of the components of a cubical three-way array T_{ijk} given in eq. (6) under the generic cyclic and acyclic permutation operators $P_{\rm cyc}$, $P_{\rm acyc}$, the Jacobi-like operator $J = T_{ijk} + T_{jki} + T_{kij}$, and the index-pair exchange of i, j denoted by P_{ij}.

Conjecture 1 *A trilinear tensor completion model which is built upon non-commutative parallel factors, i.e., that differentiates between different permutations of the same numerical values of its arguments, would in some way be "more realistic" - hence perform better than the standard CP model.*

Conjecture 1 leads to two immediate outcomes. Firstly, the standard CP model is suboptimal - since its building block is the multiplication of real arguments and is inherently commutative. If a trilinear building block is to be used, the arguments must be of dimension two at least. Likewise, the next simplest extension which is the multiplications of complex arguments, cannot be used (at least naively), as it is commutative. Secondly, in order to differentiate between all different "color" permutations of three objects, there must be at least three "colors". In other words, a single parallel factor must differentiate at least three co-clusters of each class. Non-commutative three-way relations between co-clusters must therefore involve, at the very least, a $3 \times 3 \times 3$ assignment - a mapping $\{\text{red}, \text{green}, \text{blue}\}^3 \mapsto \mathbb{R}$.

In the next Section we construct such a real trilinear approximation of three-way arrays in $\mathbb{R}^{N \times N \times N}$ for $N \geq 3$. We shall later use this construction for a general tensor completion problem.

4 The Non Commutative Latent Factors (NCLF) method

4.1 Approximating a real $N \times N \times N$ array

We recall that, given a three-dimensional cuboid array of real numbers $T \in \mathbb{R}^{N \times N \times N}$, it may be decomposed to six components according to their permutation symmetry properties. There are several options for doing this, and the decomposition we choose is

$$\begin{bmatrix} S[T] \\ A[T] \\ \mathcal{J}_{31-}[T] \\ \mathcal{J}_{31+}[T] \\ \mathcal{J}_{23-}[T] \\ \mathcal{J}_{23+}[T] \end{bmatrix}_{ijk} = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & -1 & -1 & -1 \\ 1 & 0 & -1 & 1 & 0 & -1 \\ 1 & 0 & -1 & -1 & 0 & 1 \\ 0 & 1 & -1 & 0 & -1 & 1 \\ 0 & 1 & -1 & 0 & 1 & -1 \end{bmatrix} \begin{bmatrix} T_{ijk} \\ T_{jki} \\ T_{kij} \\ T_{ikj} \\ T_{jik} \\ T_{kji} \end{bmatrix} . \tag{6}$$

Eq (6) is a list of linear combinations of T_{ijk} and its index permutations. We note that the linear mapping (6) is invertible and well-conditioned.

The symmetry properties of the six components are given in Table 1. The first two components S and A are eigenvectors of all the permutation symmetries - the first being symmetric under all permutations while the second being symmetric under cyclic (even) permutations and anti-symmetric under acyclic (odd) ones. The next four components are eigenvectors of only a single permutation symmetry each, but all satisfy a Jacobi-like identity:

$$\mathcal{J}[T] := T_{ijk} + T_{jki} + T_{kij} \equiv 0. \tag{7}$$

We use the images of these operators to define three linear subspaces of $\mathbb{R}^{N\times N\times N}$. The first two are the images of the totally symmetric and totally antisymmetric operators $Im(S)$ and $Im(A)$. The third subspace is the sum of the images of the last four operators, which is also equal to the kernel of the Jacobi identity $Im\mathcal{J}_{23+} + Im\mathcal{J}_{23-} + Im\mathcal{J}_{31+} + Im\mathcal{J}_{31-} = Ker(\mathcal{J})$. Direct calculation gives that, taken as subspaces of $\mathbb{R}^{N\times N\times N}$ with the Euclidean inner product associated with the Froebenius norm, the three spaces are pairwise orthogonal and span the full space, hence $\mathbb{R}^{N\times N\times N} = Im(A) \oplus Im(S) \oplus Ker(J)$.

Next, we construct diagonal trilinear approximations of for each of these six components, which satisfy the relevant symmetries. The second component $A[T]$ is approximated using the standard totally antisymmetric form, or standard triple product in $\mathbb{R}^3$, which is equal to $\det[u\ v\ w] = u(v \times w)$, with three-dimensional latent factors $u, v, w \in \mathbb{R}^3$. In the next two Sections, we approximate the other five components using a two-step process:

1. In Section 4.2 we define a trilinear non-commutative mapping, which we shall denote by μ, over a two-dimensional subspace of $\mathbb{R}^{2\times 2}$. As it is two dimensional, it is hard to think of a simpler such mapping.

2. Next, in Section 4.3 we apply the symmetrizing operators of (6) on this trilinear form μ, to obtain the approximations for the five components.

In Section 5 we provide numerical indications that each of these two steps improves the overall approximation of the chosen datasets.

4.2 The space $\mathbb{C}^\perp$ and operation μ

Let us look for the simplest "atom" for the Jacobi components - that is the simplest possible space supporting a noncommutative trilinear product. This space is the key component of our mathematical model. We note that the complex version of this space has been used in computational Physics of three-color quantum models [Ker10].

A trilinear operation with one dimensional real arguments must be commutative, and so such a space must have at least two dimensional arguments. Non-commutativity and trilinearity leads us towards 2×2 matrix multiplication as a representation.

Before we continue, let us recall two basic facts on the space of 2×2 real matrices $\mathbb{R}^{2\times 2}$. First, it is spanned by the identity matrix and the three Pauli spin matrices:

$$\mathbf{I} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \qquad \sigma_1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \qquad i\sigma_2 = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}, \qquad \sigma_3 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix},$$

which are mutually orthogonal in the inner product associated with the Froebenius norm. In other words they are an orthogonal basis of $\mathbb{R}^{2\times 2}$. Second, the space of complex numbers $\mathbb{C}$ is isomorphic, using the Cayley-Dickson construction, to the space of antisymmetric 2×2 real matrices of the form

$$\mathrm{CD}(\mathbb{C}) = \left\{ \left[\begin{array}{cc} c^0 & c^2 \\ -c^2 & c^0 \end{array} \right] \middle| c^0, c^2 \in \mathbb{R} \right\} = \left\{ c^0\mathbf{I} + c^2 i\sigma_2 \middle| c^0, c^2 \in \mathbb{R} \right\}$$

with matrix multiplication corresponding to the product of complex numbers. In this subspace of $\mathbb{R}^{2\times 2}$, matrix multiplication is commutative.

With these facts in mind, we therefore turn to the orthogonal complement $\mathbb{C}^\perp$ of $\mathrm{CD}(\mathbb{C})$ to look for non-commutative trilinear operations. From the fact that $\{\mathbf{I}, \sigma_1, \sigma_2, \sigma_3\}$ is an orthogonal basis it immediately follows that $\mathbb{C}^\perp$ is the span of $\{\sigma_1, \sigma_3\}$:

$$\mathbb{C}^\perp := \left\{ c^1\sigma_1 + c^3\sigma_3 \middle| c^1, c^3 \in \mathbb{R} \right\} = \left\{ \left[\begin{array}{cc} c^3 & c^1 \\ c^1 & -c^3 \end{array} \right] \middle| c_1, c_3 \in \mathbb{R} \right\}. \tag{8}$$

It is also the space of traceless symmetric 2×2 real matrices.

Additionally, for each ordered triplet $u, v, w \in \mathbb{C}^\perp$, setting

$$u = u^1\sigma_1 + u^3\sigma_3, \tag{9}$$

and similarly for v, w, direct calculation shows that $\mathbb{C}^\perp$ is closed under a triple matrix product:

$$\begin{aligned} uvw &= \left(u^1v^1w^1 + u^3v^3w^1 - u^3v^1w^3 + u^1v^3w^3\right)\sigma_1 \\ &\quad + \left(u^3v^3w^3 + u^1v^1w^3 - u^1v^3w^1 + u^3v^1w^1\right)\sigma_3. \end{aligned}$$

Hence, the mapping

$$\begin{array}{rcl} \mu : \mathbb{C}^\perp \times \mathbb{C}^\perp \times \mathbb{C}^\perp & \to & \mathbb{C}^\perp \\ \mu(u, v, w) & \mapsto & uvw \end{array} \tag{10}$$

is a well defined real trilinear operation. Considering commutativity, the product uvw is symmetric with respect to exchange of the first and third parameters, but not to a permutation which changes the second argument[5]

$$uvw = wvu, \qquad uvw \neq uwv, \qquad u, v, w \in \mathbb{C}^\perp. \tag{11}$$

We note that $\mathbb{C}^\perp$ is not closed under the standard (binary) matrix multiplication - for $u, v \in \mathbb{C}^\perp$ we have $uv \in CD(\mathbb{C})$, not $\mathbb{C}^\perp$. Therefore, $\mathbb{C}^\perp$ is not a group under matrix multiplication, and is hence not an algebra, but rather a ternary algebra. Similarly to the standard triple product in $\mathbb{R}^3$, the pair $(\mathbb{C}^\perp, \mu)$ is a purely third-order construct.

[5] Indeed, the algebra $Cl(1,1)$ is defined as the two dimensional Clifford Algebra having one symmetric and one antisymmetric index.

4.3 Approximating the five components

Here, we approximate the symmetric and Jacobi components of T_{ijk}, which are S and $\mathcal{J}_{31\pm}, \mathcal{J}_{23\pm}$, using linear combinations of the form μ on $\mathbb{C}^{\perp}$. Specifically, if the latent factor corresponding to item i is

$$u_i = u_i^1 \sigma_1 + u_i^3 \sigma_3 \in \mathbb{C}^{\perp},$$

and similarly for v_j and w_k, we apply the symmetrizing operators of eq. (6) on $\mu(u_i, v_j, w_k)$ to obtain these operators as explicit cubic polynomials of the coefficients. For example, the totally symmetric component is

$$\begin{aligned} S(u_i, v_j, w_k) &:= u_i v_j w_k + v_j w_k u_i + w_k u_i v_j + u_i w_k v_j + w_k v_j u_i + v_j u_i v_j \\ &= 2\left(3u_i^1 v_j^1 w_k^1 + u_i^3 v_j^1 w_k^1 + u_i^1 v_j^3 w_k^1 + u_i^1 v_j^1 w_k^3\right)\sigma_1 \\ &\quad +2\left(3u_i^3 v_j^3 w_k^3 + u_i^1 v_j^3 w_k^3 + u_i^3 v_j^1 w_k^3 + u_i^3 v_j^3 w_k^1\right)\sigma_3, \end{aligned} \tag{12a}$$

and similarly

$$\begin{aligned} \mathcal{J}_{31-}(u_i, v_j, w_k) &= 2\left(u_i^1 v_j^3 w_k^3 - u_i^3 v_j^3 w_k^1\right)\sigma_1 \\ &\quad +2\left(u_i^3 v_j^1 w_k^1 - u_i^1 v_j^1 w_k^3\right)\sigma_3 \end{aligned} \tag{12b}$$

$$\begin{aligned} \mathcal{J}_{31+}(u_i, v_j, w_k) &= 2\left(u_i^1 v_j^3 w_k^3 - 2u_i^3 v_j^1 w_k^3 + u_i^3 v_j^3 w_k^1\right)\sigma_1 \\ &\quad +2\left(u_i^3 v_j^1 w_k^1 - 2u_i^1 v_j^3 w_k^1 + u_i^1 v_j^1 w_k^3\right)\sigma_3 \end{aligned} \tag{12c}$$

$$\begin{aligned} \mathcal{J}_{23-}(u_i, v_j, w_k) &= 2\left(u_i^3 v_j^3 w_k^1 - u_i^3 v_j^1 w_k^3\right)\sigma_1 \\ &\quad +2\left(u_i^1 v_j^1 w_k^3 - u_i^1 v_j^3 w_k^1\right)\sigma_3 \end{aligned} \tag{12d}$$

$$\begin{aligned} \mathcal{J}_{23+}(u_i, v_j, w_k) &= 2\left(-2u_i^1 v_j^3 w_k^3 + u_i^3 v_j^1 w_k^3 + u_i^3 v_j^3 w_k^1\right)\sigma_1 \\ &\quad +2\left(-2u_i^3 v_j^1 w_k^1 + u_i^1 v_j^3 w_k^1 + u_i^1 v_j^1 w_k^3\right)\sigma_3. \end{aligned} \tag{12e}$$

Importantly, the symmetry (11) of μ implies that the completely anti-symmetric combination vanishes

$$A(u_i, v_j, w_k) := u_i v_j w_k + v_j w_k u_i + w_k u_i v_j - (w_k v_j u_i + u_i w_k v_j + v_j u_i w_k) \equiv 0.$$

This is reassuring, as the Jacobi and symmetric components are orthogonal to the anti-symmetric component.

4.4 The general cuboid case

The previous subsections dealt with a cubical array in $\mathbb{R}^{N\times N\times N}, N \geq 3$. We shall reuse the same model in the general cuboid case as is, without any formal justification. The intuition behind this is that the previous derivation applies to modeling the relations of co-clusters (aka "colors"), which can be cubical even if the approximated tensor is a cuboid. The ultimate judge is, of course, empirical evidence.

Therefore, combining the results of this Section, given a three-dimensional (cuboid) array of real num-

bers $T \in \mathbb{R}^{I\times J\times K}$, we approximate it as

$$\begin{aligned} T^{NCLF}_{ijk} &= b_0 + b_{1i} + b_{2j} + b_{3k} + \sum_{r=1}^{R_s} \zeta^S_r S\left(u^S_{ir}, v^S_{jr}, w^S_{kr}\right) + \sum_{r=1}^{R_A} \alpha_r \det\left[u^A_{ir}, v^A_{jr}, w^A_{kr}\right] \\ &+ \sum_{p\in 23,31} \sum_{s=\pm} \sum_{r=1}^{R_{ps}} \zeta^{(ps)}_r \mathcal{J}_{ps}\left(u^{(ps)}_{ir}, v^{(ps)}_{jr}, w^{(ps)}_{kr}\right), \\ &(\cdot)^S,\ (\cdot)^{(ps)} \in \mathbb{C}^{\perp}, \quad (\cdot)^A \in \mathbb{R}^3, \end{aligned} \tag{13}$$

where $b_{(\cdot)}$ are corresponding bias terms, the operators $S, \mathcal{J}_{ps}$ are as defined in (12), $\det[\cdot]$ is the standard triple product in $\mathbb{R}^3$ and the quantities $\zeta^{(\cdot)}_r \in \mathbb{R}^2$, which generalize singular values, imply summation over the σ_1 and σ_3 components.

Equation (13) is the concrete, explicit model of the general form (3), and is the key result of this paper. Note that this approximation is as close to diagonal as possible, while still being noncommutative, i.e., while differentiating between different permutations of the latent factors, as required by Conjecture 1.

5 Numerical Experiments

Here we present the results of numerical experiments for two public datasets - the MovieLens Dataset [Gro14] and the Fannie Mae Single-Family Loan Performance dataset [Mae14]. The goal of experiments was a comparison of the expansion (13) with the standard CP model, rather than obtaining the optimal model for each Dataset. In both cases we used a binary response variable and a logistic-regression model, so that the probability of a positive event is modeled by (4) and training consists of solving the minimization problem (5), see Section 2.

5.1 Benchmark Approximations

Five benchmark approximations of the logodds T_{ijk} were compared:

1. A bias-only method, which is equivalent to a Naive Bayes approximation:

$$T^{\text{Bias only}}_{ijk} = b_0 + b_{1i} + b_{2j} + b_{3k}.$$

   The total logodds bias b_0 and the relative biases b_{fi} for each entity i of factor $f = 1, 2, 3$ were estimated as empirical logodds

$$b_0 = \log\frac{P+1}{N+1}, \qquad b_{fi} = \log\frac{P_{fi}+1}{N_{fi}+1} - b_0. \tag{14}$$

   where P, N are the total counts of positive and negative events for the training set and P_{fi}, N_{fi} are the same counts for each entity fi.

2. The standard CP approximation (1) with a latent dimension equal to that of the NCLF method $R = 13$.

3. The standard CP with the best latent dimensions $R = 5$ for both the MovieLens and Fannie Mae datasets. The best dimensions were chosen via nine-fold cross-validation.

4. In order to test the utility of the derivation of Section 4.3, i.e., of using the separate approximations (12) for each of the five components S and $\mathcal{J}_{(\cdot)}$, we also benchmark a "primitive" NCLF approximation given by

$$T_{ijk}^{\text{primitive NCLF}} = b_0 + b_{1i} + b_{2j} + b_{3k} + \det\left[u_{ir}^A, v_{jr}^A, w_{kr}^A\right] + \sum_{r=1}^{5} \zeta_r \mu\left(u_{ri}, v_{rj}, w_{rk}\right). \tag{15}$$

   This approximation explicitly models only the totally-antisymmetric component A, while using the primitive operation μ instead of modeling each of the five components S and $\mathcal{J}_{(\cdot)}$. We recall that μ has partial symmetry (11). This implies that the partially-antisymmetric components $\mathcal{J}_{23-}, \mathcal{J}_{31-}$ are not approximated by (15), while the rest of the components are.

5. The proposed NCLF method, wherein T_{ijk} is given by (13), and each of the components has a single latent factor $R_{(\cdot)} = 1$.

Models were trained using the Stochastic Gradient Descent method (SGD) of the momentum variant, with decreasing time-steps. In all the approximations 1-5, the bias terms were taken to be identical. Specifically, they were not trained by SGD but rather chosen, before the SGD simulations, by (14). The parallel factors were regularized using the L^2 norm, using nine-fold cross-validation to pick the regularization parameter, and 25-fold cross-validation to measure performance of the best configuration.

5.2 The Datasets

The MovieLens Dataset [Gro14] contains a million user-ratings of movies on a scale of one to five. Ratings of 4 and 5 were considered to be positive events, and lower ratings as negative events. Overall, 424928 negative and 575281 positive rating events were considered. The three factors we consider are those of item, user and hour of week (totaling 168 bins).

The Fannie Mae Single-Family Loan Performance dataset [Mae14] is a publicly available dataset which, at the time of submission, holds fixed rate prime mortgage acquisition and performance data, at monthly resolution, for the period from January 1999 till June 2013, including. Only first-time home buyers whose loan purchase was buying or undefined were considered. The three factors chosen where credit-score, property location denoted by property state and 3-digit zip code, and origination month. We chose not to group or smooth different values of credit scores or time periods longer than a month, so as not to make the prediction problem easier. A mortgage was considered to have defaulted if delinquent more than 150 days over the full period. Non-default events were uniformly downsampled. Overall, 1197549 non-default and 876707 default acquisition events were considered.

5.3 Results

Cross-validation performance of the five approximations of 5.1 applied to the MovieLens dataset is given in Table 2, and their performance over the Fannie Mae dataset is given in Table 3. In both cases, we see that the NCLF models considerably outperforms the standard CP model of the same latent dimension 13, and significantly outperforms CP models of lower dimensions, as measured by all metrics: AUC, L_1 error and L_2 error.

MovieLens							
Method		AUC	ΔAUC	L1	ΔL1	L2	ΔL2
1	Bias only	0.6494	21	0.4542	6	0.4712	5
2	CP, $R = 13$	0.7625	36	0.3387	21	0.4456	22
3	best CP, $R = 5$	0.7783	56	0.3470	31	0.4318	26
4	primitive NCLF	0.7817	48	0.3536	39	0.4283	25
5	NCLF	0.7920	31	0.3365	23	0.4256	18
	NCLF-best CP	0.0137	45	0.0105	27	0.0062	22

Table 2: Performance of the five approximations as given in Section 5.1, for the MovieLens 1M ratings dataset, obtained by 25-fold cross-validation. Columns denoted by $\Delta(\cdot)$ give sample standard errors, multiplied by 10^4. The last row gives the absolute difference of the CP with the best rank $R = 5$ to the NCLF.

Fannie Mae							
	Method	AUC	ΔAUC	L1	ΔL1	L2	ΔL2
1	Bias only	0.7689	58	0.3819	7	0.4329	6
2	CP, $R = 13$	0.8242	38	0.3028	14	0.4105	25
3	best CP, $R = 5$	0.8326	55	0.3062	41	0.4029	33
4	primitive NCLF	0.8447	19	0.3040	12	0.3954	12
5	NCLF	0.8462	16	0.3029	9	0.3942	10
	NCLF-best CP	0.0136	41	0.0033	30	0.0087	24

Table 3: Same as Table 2, for the Fannie Mae dataset.

The numerical experiments therefore strongly corroborate Conjecture 1, at least for these datasets and with the SGD numerical method - under these assumptions, non-commutative latent factors outperform the standard CP.

Additionally, there is weak evidence that the proposed NCLF mildly outperforms the "primitive NCLF" model (15), meaning that applying the symmetrizing operators of Section 4.3 (thereby approximating the two components $\mathcal{J}_{23-}, \mathcal{J}_{31-}$) provides an improved approximation.

6 Discussion and Future Directions

In this study, we develop a novel tensor-completion method for three-way arrays, which is both diagonal and built upon non-commutative latent factors. In order to do this, we apply symmetrizing operations on the simplest non-commutative purely trilinear operation we could find - that of three-matrix product on a two-dimensional space. We test our model and numerical method on a binary-response supervised-learning problem from two publicly-available datasets, finding that it outperforms the CP model.

The specific application we are interested in is modeling sparse, large-scale three-way relations in the supervised-learning setting, i.e., in three-way CF problems. However, we find no apriori reason that this model may not be extended to a broader setting. Some future avenues for research include:

1. **Unsupervised learning:** An interesting question is if and how much a non-commutative model may be used to discover non-commutative patterns in three-way-relation data. The intuitions leading to its development in Section 3 should still apply.

2. **(Dense) Tensor Factorization:** A possible future direction may be the analysis of this model in the context of tensor-factorization - i.e., of approximation a full tensor with no missing values. We note that in this setting there are Fourier-based generalizations of the SVD [KM11] in addition to the HOSVD of Delathauwer et al., and a comparison of the three options may be interesting.

3. **Extension to Quaternions:** The space $(\mathbb{C}^{\perp}, \mu)$ is in fact a two-dimensional subspace of the ring of quaternions. One may consider applying the symmetrizing operators (6) on three-quaternion products instead of on μ - in fact, this was the original direction of this work. The resulting approximation might be more expressive than NCLF, but have a double latent dimension, and so be more likely to overfit. Nevertheless, in a world where the volume of data keeps increasing, such an extension might some day prove superior.

In summary, Non Commuting Latent Factors present a simple, scalable extension of the CP model which outperforms it on the two datasets tried.

References

[CLdA09] P. Comon, X. Luciani, and A. L. F. de Almeida. Tensor decompositions, alternating least squares and other tales. *Journal of Chemometrics*, 3:393–405, Aug. 2009.

[dLdMV00] L. de Lathauwer, B. de Moor, and J. Vandewalle. A multilinear singular value decomposition. *SIAM J. Matrix Anal. Appl.*, 21:1253–1278, 2000.

[dSL08] V. de Silva and L.-H. Lim. Tensor rank and the ill-posedness of the best low-rank approximation problem. *SIAM J. Matrix Anal. Appl.*, 30(3):1084–1127, 2008.

[GKT13] L. Grasedyck, Daniel Kressner, and Christine Tobler. A literature survey of low-rank tensor approximation techniques. (1302.7121), 2013.

[Gro14] GroupLens. MovieLens dataset, 2014.

[Hit27a] F. L. Hitchcock. The expression of a tensor or a polyadic as a sum of products. *J. Math. Phys*, 6(1):164–189, 1927.

[Hit27b] F. L. Hitchcock. Multiple invariants and generalized rank of a p-way matrix or tensor. *J. Math. Phys*, 7(1):39–79, 1927.

[KB09] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.

[Ker10] Richard Kerner. Cubic and ternary algebras, ternary symmetries, and the lorentz group. In *RIMS Kokyuroku published in 2014, No. 1872-1920*, volume 1705, pages 134–146, 2010.

[KM11] Misha E. Kilmer and Carla D. Martin. Factorization strategies for third-order tensors. *Linear Algebra Appl.*, 435(3):641–658, 2011.

[Kol01] Tamara G Kolda. Orthogonal tensor decompositions. *SIAM J. Matrix Anal. Appl.*, 23(1):243–255, 2001.

[KSV] Daniel Kressner, Michael Steinlechner, and Bart Vandereycken. Low-rank tensor completion by riemannian optimization. *BIT Numerical Mathematics*, pages 1–22.

[Mae14] Fannie Mae. Fannie mae single-family loan performance data, 2014.

[Tuc63] L.R. Tucker. Implications of factor analysis of three-way matrices for measurement of change. *Problems in measuring change*, pages 122–137, 1963.

[Tuc64] L. R. Tucker. The extension of factor analysis to three-dimensional matrices. *Contributions to mathematical psychology*, pages 109–127, 1964.

[Tuc66] L. R. Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31:279–311, 1966.